

WIICA


Workload ISA-Independent Characterization for Applications

Yakun Sophia Shao, Emma Wang,
Gu-Yeon Wei, David Brooks
Harvard University




Specialized architectures are decoupled from legacy ISAs.

Spectrum of Specialization:



Specialized architectures are decoupled from legacy ISAs.

Spectrum of Specialization:



Performance-Counter Based Workload Characterization

- Metrics
 - IPC
 - Cache miss rates
 - Branch mis-prediction rates
 - ...
- Microarchitecture-dependent
 - What if there is a bigger cache/a better branch predictor?
 - Not program intrinsic characteristics

ISA impacts program behaviors.

Stack Overhead

- Limited Registers
- Additional Load/Store

ISA impacts program behaviors.

Stack Overhead

- Limited Registers
- Additional Load/Store

Complex Operations

- Memory Operands
- Vector Operations

ISA impacts program behaviors.

Stack Overhead

- Limited Registers
- Additional Load/Store

Complex Operations

- Memory Operands
- Vector Operations

Calling Conventions

WIICA Summary

Goal:

- An analysis tool to characterize workloads ISA-independent characteristics for specialized architectures

WIICA Summary

Goal:

- An analysis tool to characterize workloads ISA-Independent characteristics for specialized architectures

Methods:

- Leverage compiler's intermediate representation (IR)
- Categorize characteristics into compute, memory, and control

WIICA Summary

Goal:

- An analysis tool to characterize workloads ISA-Independent characteristics for specialized architectures


Methods:

- Leverage compiler's intermediate representation (IR)
- Categorize characteristics into compute, memory, and control


Takeaways:

- ISA-dependent characterization is misleading for specialization.
- ISA-independent characterization allows designers to quickly identify opportunities for specialization.

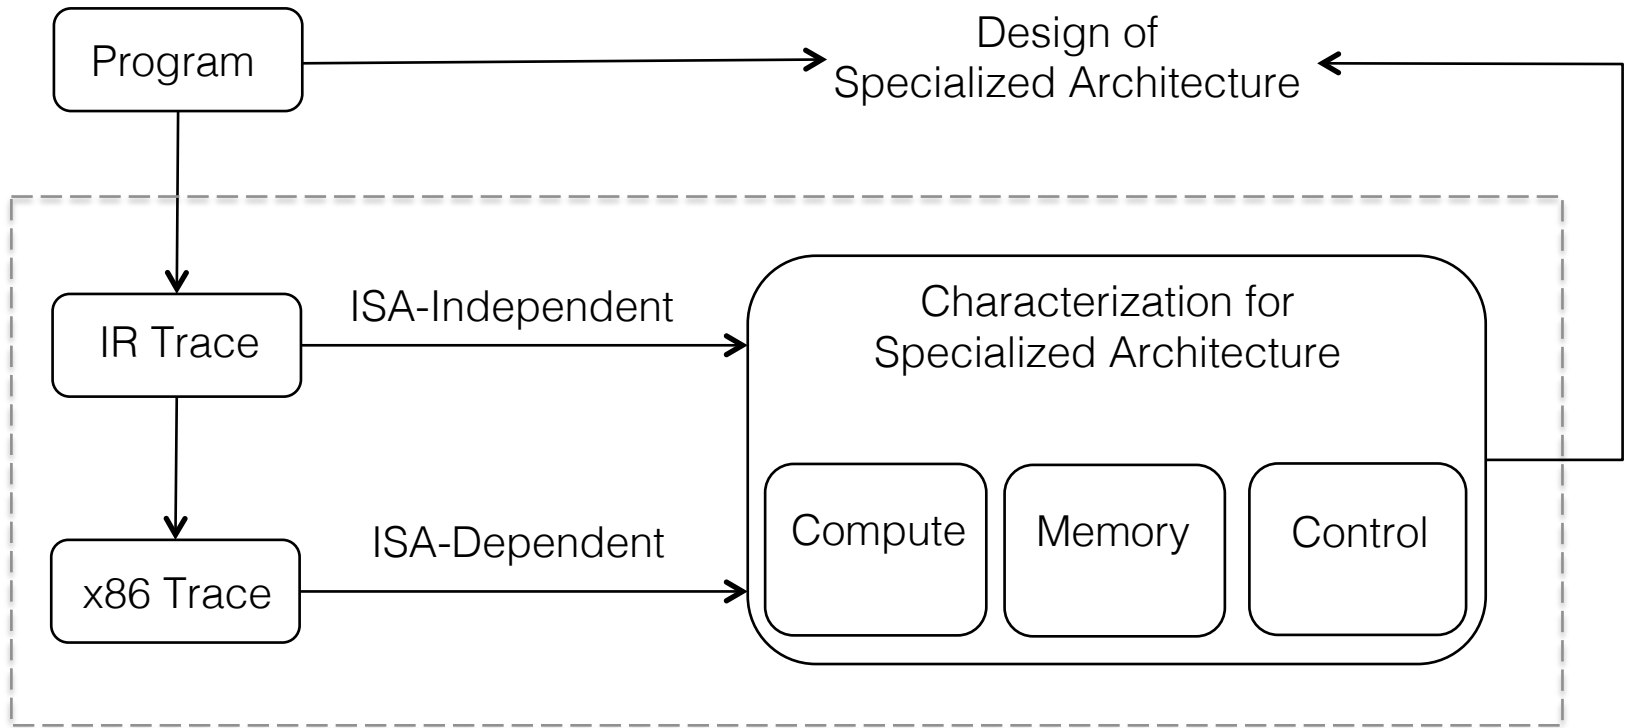
Tool Overview




Tool Overview




Tool Overview



Program Representations



Program Representations




ILDJIT

- A modular compilation framework
- Performs machine-independent classical optimizations at the IR level
- Uses LLVM's back end to
 - Do machine-dependent optimizations
 - Generate machine code

Campanoni, et al., A Highly Flexible, Parallel Virtual Machine: Design and Experience of ILDJIT, Software Practice Experience, 2010


Program Representations



ILDJIT IR

- High-level IR
- Machine-, ISA-, and system-library-independent
- Features:
 - 80 instructions
 - Unlimited registers
 - Only loads/stores access memory
 - No vector operations
 - Parameters are passed by variables


Program Representations



x86 Trace

- Used for ISA-dependent analysis
- Semantically equivalent to the IR code
- Collected with Pin instrumentation

Tool Overview



ISA-Independent Workload Characteristics

Compute

- Opcode Diversity
- Static Instructions (I-MEM)

Memory

- Memory Footprint (D-MEM)
- Memory Entropy
- Locality Score

Control

- Branch Instruction Counts
- Branch Entropy

ISA-Independent Workload Characteristics

Compute

- Opcode Diversity
- **Static Instructions (I-MEM)**


Memory

- Memory Footprint (D-MEM)
- Memory Entropy
- Locality Score

Control

- Branch Instruction Counts
- Branch Entropy

Compute::Static Instructions



ISA-Independent Workload Characteristics

Compute

- Opcode Diversity
- Static Instructions (I-MEM)

Memory

- Memory Footprint (D-MEM)
- **Memory Entropy**
- Locality Score

Control

- Branch Instruction Counts
- Branch Entropy

Memory::Entropy

Entropy: a measure of the randomness

$$Entropy = - \sum_{i=1}^N p(x_i) * \log_2 p(x_i)$$

Memory::Entropy

Entropy: a measure of the randomness

$$\text{Entropy} = - \sum_{i=1}^N p(x_i) * \log_2 p(x_i)$$

Case 1:

X is always a constant.

$$p(X) = 1$$

$$\log_2 p(X) = 0$$

$$\text{Entropy} = 0$$

Case 2:

N possible outcomes of X occur equally.

$$p(X) = \frac{1}{N}$$

$$\log_2 p(X) = \log_2 N^{-1}$$

$$\text{Entropy} = -N * \frac{1}{N} * \log_2 N^{-1}$$

$$\text{Entropy} = \log_2 N$$



Memory::Global Address Entropy

Temporal Locality

Address Stream A
(less temporal locality)

0000
0001
0010
0011

Entropy = 2

Address Stream B
(more temporal locality)

0011
0011
0011
0011

Entropy = 0

Yen, Draper, and Hill. Notary: Hardware Techniques to Enhance Signatures. MICRO 08

Memory::Global Address Entropy

Temporal Locality

Address Stream A
(less temporal locality)


0000
0001
0010
0011

Entropy = 2

Address Stream B
(more temporal locality)

0011
0011
0011
0011

Entropy = 0



Yen, Draper, and Hill. Notary: Hardware Techniques to Enhance Signatures. MICRO 08



Memory::Global Address Entropy

Temporal Locality

Address Stream A
(less temporal locality)


0000
0001
0010
0011

Entropy = 2

Address Stream B
(more temporal locality)

0011
0011
0011
0011

Entropy = 0

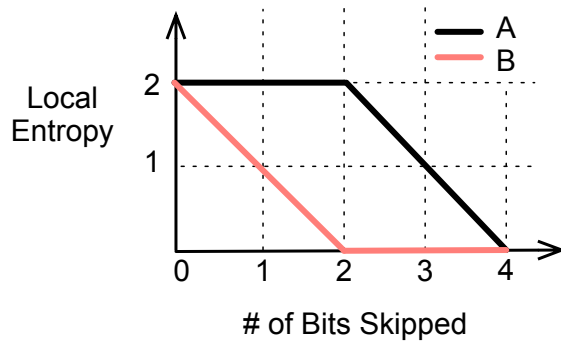
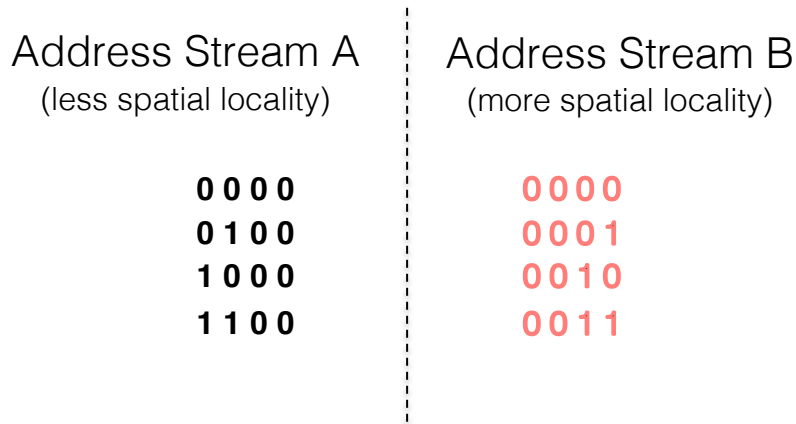


Yen, Draper, and Hill. Notary: Hardware Techniques to Enhance Signatures. MICRO 08




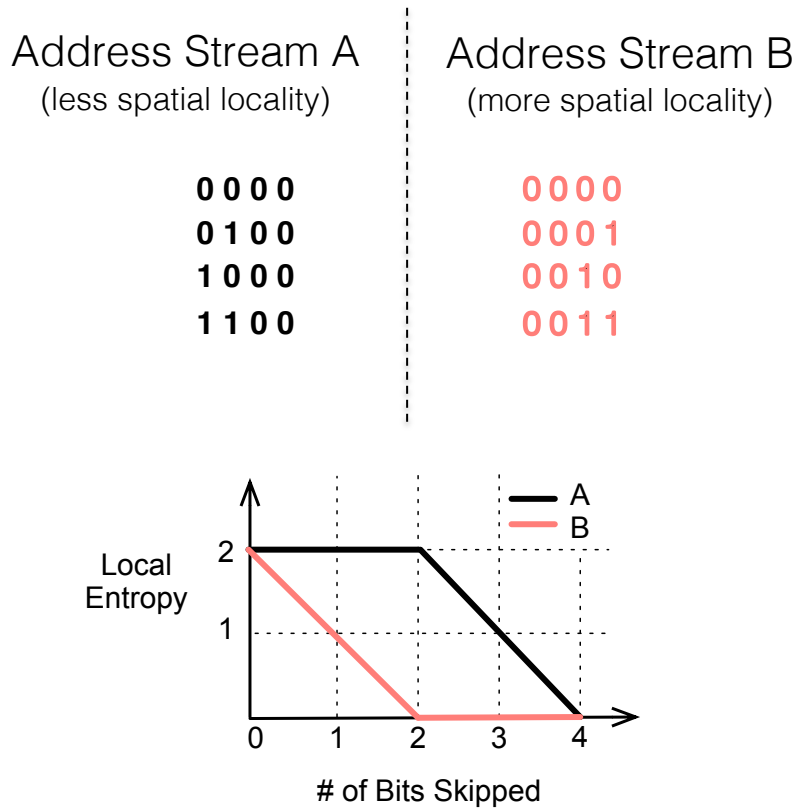
Memory::Local Address Entropy

Spatial Locality



Memory::Local Address Entropy

Spatial Locality



Memory::Spatial Locality Score

Address Stream A
(less spatial locality)

0x0000
0x0004
0x0008
0x000C

$$L = \frac{P(\text{stride} = 4)}{\text{stride} = 4}$$
$$= \frac{1}{4}$$

Address Stream B
(more spatial locality)

0x0000
0x0001
0x0002
0x0003

$$L = \frac{P(\text{stride} = 1)}{\text{stride} = 1}$$
$$= 1$$

$$L_{\text{spatial}} = \sum_{\text{stride}=1}^{\text{stride}<\infty} \frac{P(\text{stride})}{\text{stride}}$$

Weinberg, et al, *Quantifying Locality in the memory access patterns of HPC applications*,
In SC, 2005



ISA-Independent Workload Characteristics

Compute

- Opcode Diversity
- Static Instructions (I-MEM)

Memory


- Memory Footprint (D-MEM)
- Memory Entropy
- Locality Score

Control

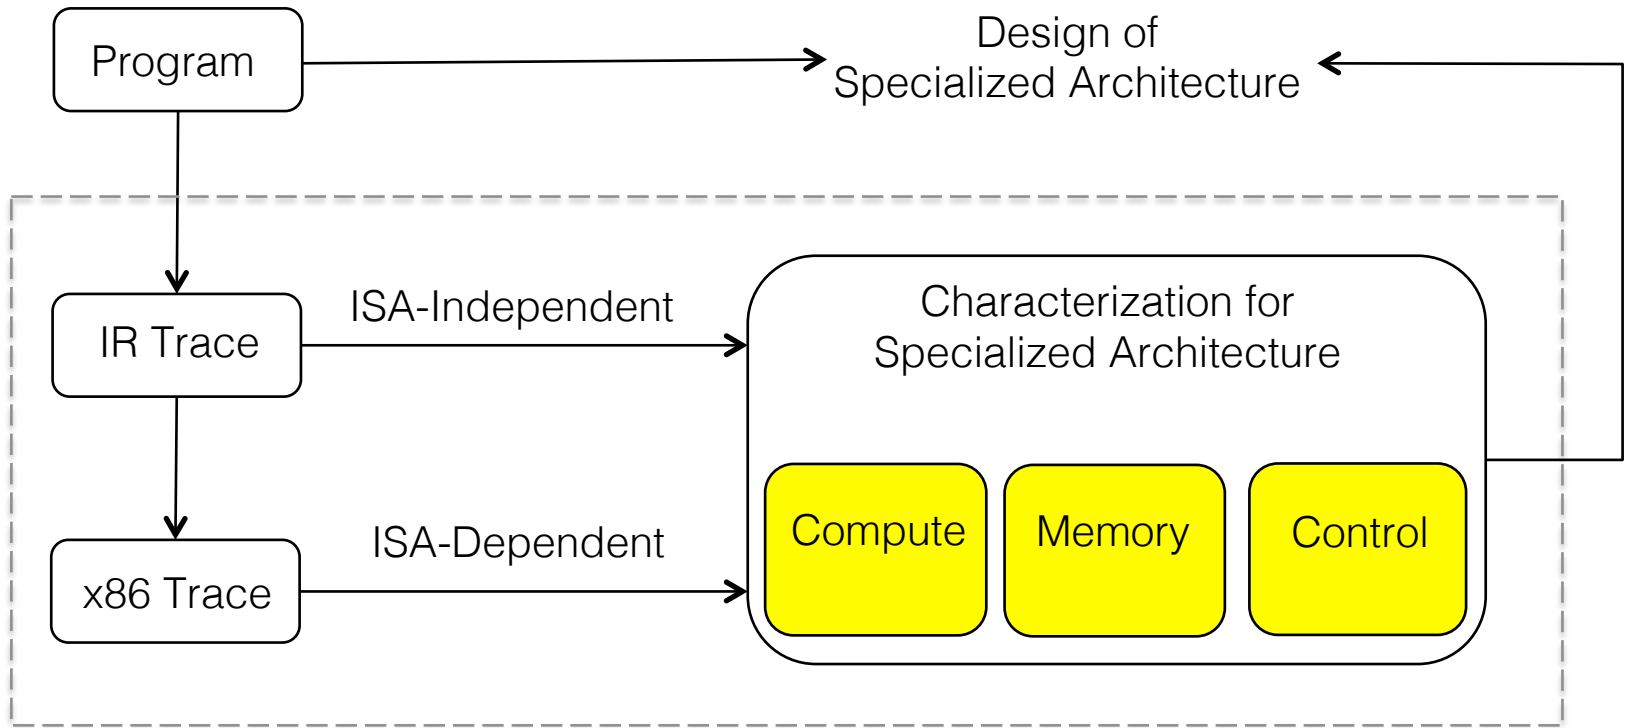
- Branch Instruction Counts
- **Branch Entropy**

Yokota, et al, Introducing Entropies for Representing Program Behavior and Branch Predictor Performance, 07

Control::Branch Entropy



Tool Overview



Publications

- Implications of the Power Wall: Dim Cores and Reconfigurable Logic
 - Wang and Skadron, IEEE Micro, 2013
- ISA-Independent Workload Characterization and its Implications for Specialized Architectures
 - Shao and Brooks, ISPASS, 2013